

We claim:

1. A multi-stage classification method of matching a plurality of strings to one or more policies embodied in a plurality of stages, each stage comprising one or more entries, each entry comprising a criterion and a pointer to a subsequent stage; the method
5 comprising the steps of:

(a) generating a current pointer to a current stage using a current string;
(b) searching one or more criteria in the current stage, the searching step further comprising the steps of:

(i) comparing the current string to the criteria of one or more entries
10 beginning with entries indicated by the one or more current pointers,
(ii) defining the one or more current pointers equal to one or more new pointers, wherein each of the new pointers is associated with a criterion in the current stage that matches the current string, and
(iii) incrementing the current stage and the current string;

15 (c) repeating the searching step for each of the plurality of stages;
(d) generating a score for each match identified in each of the plurality of stages;
and

(e) selecting the policy of the one or more policies associated with a highest cumulative score, the cumulative score being the sum of the score for each
20 match at each stage of the plurality of stages.

2. The multi-stage classification method of claim 1, wherein each of the plurality of stages comprises a hierarchical table, wherein each hierarchical table comprising
25 criteria associated with the one or more policies.

3. The multi-stage classification method of claim 2, wherein the criteria are arranged in one or more criterion groups, wherein each criterion of a criterion group in a child hierarchical table shares a common criterion in a parent hierarchical table.

4. The multi-stage classification method of claim 3, wherein the step of comparing the current string to the criteria of one or more entries is limited to the criteria of one or more criterion groups beginning with entries indicated by the one or more current pointers.
5. The multi-stage classification method of claim 1, wherein the one or more criteria comprise one or more numerical values.
6. The multi-stage classification method of claim 5, wherein the one or more criteria comprise one or more wildcard operators.
7. The multi-stage classification method of claim 5, wherein the one or more criteria comprise one or more logical operations.
8. The multi-stage classification method of claim 7, wherein the one or more logical operations are selected from the group consisting of: Boolean operations, relational operations, string operations, set operations, and a combination thereof.
9. The multi-stage classification method of claim 6, wherein wildcard operators are assigned a lower score than the score assigned to an exact match with a numerical value.
10. The multi-stage classification method of claim 1, wherein the current pointer of the current pointer generating step is generated by hashing the first string of the plurality of strings.
11. The multi-stage classification method of claim 1, wherein the plurality of strings comprise a plurality of nibbles derived from a protocol data unit (PDU).

12. The multi-stage classification method of claim 11, wherein the plurality of nibbles are selected from one or more fields of the PDU, the fields being selected from the group consisting of: source address, destination address, input port on which the PDU was received, source transmission control protocol (TCP) port, source user datagram port (UDP), destination UDP, layer 3 packet type, layer 4 packet type, TCP acknowledgement flag, acknowledgement bit, type of service value, and PDU payload.
13. The multi-stage classification method of claim 3, wherein the step of comparing comprises comparing the current string to the criteria of one or more entries until a match is detected.
14. The multi-stage classification method of claim 13, wherein the method further includes a reverting step, subsequent to the searching step and prior to the repeating step, comprising:
- (a) decrementing the current stage and the current string, whereby searching may return to a preceding criterion group of a preceding hierarchical table,
 - (b) selecting a criteria of said preceding criterion group not previously compared with the current string, and
 - (c) resuming the searching step.
15. The multi-stage classification method of claim 13, wherein the reverting step further includes the step of incrementing a depth counter.
16. The multi-stage classification method of claim 14, wherein the method terminates if the depth counter has been incremented a predetermined number of times.
17. The multi-stage classification method of claim 13, wherein the searching step further includes the step of:
- (a) generating an estimated cumulative score at an intermediate hierarchical table,

- (b) determining whether the estimated cumulative score could satisfy a predetermined score threshold, and
- (c) ending the searching step if the estimated cumulative score fails could not satisfy the score threshold.

5

18. The multi-stage classification method of claim 1, wherein the number of stages is greater than two.

10 19. The multi-stage classification method of claim 1, wherein the pointer in which the last hierarchical table associated with the selected policy points to an action in a rules table.

15 20. The multi-stage classification method of claim 11, wherein one or more of the plurality of nibbles is extracted from a PDU Internet Protocol (IP) version 4 or IP version 6 address.

20 21. A multi-stage classification method of selecting a policy from a plurality of policies based upon a plurality of protocol data unit (PDU) nibbles, the plurality of policies being embodied in a plurality of hierarchical tables, the plurality of hierarchical tables comprising a first table and one or more child tables, the method comprising the steps of:

- (a) generating, from a first nibble, a first pointer to a first table of a plurality of tables that embody the policies, each table comprising one or more criteria and associated pointer;
- (b) comparing the first nibble with each of the one or more criteria of the first table beginning with the criterion indicated by the first pointer;
- (c) identifying one or more matches between the first nibble and one or more associated pointers to a child table;

- (d) searching a current child table with a current nibble from the plurality of nibbles, the searching step further comprising:
- 5 (i) comparing the current nibble with one or more criteria of the child table, beginning with the criteria indicated by each of the one or more pointers obtained in a previous identifying step, and
- (ii) identifying one or more matches between the current nibble and one or more pointers to a next child table, each pointer associated with a match;
- (e) repeating the searching step for each child table of the plurality of tables;
- (f) determining a score for each match identified in each hierarchical table; and
- 10 (g) selecting the policy of the one or more policies associated with a highest cumulative score, the cumulative score being the sum of the scores for each match at each of the plurality of hierarchical tables.
22. The multi-stage classification method of claim 21, wherein the criteria of each
- 15 hierarchical table are arranged in one or more criterion groups, wherein each criterion of a criterion group in a child hierarchical table shares a common criterion in an associated parent hierarchical table.
23. The multi-stage classification method of claim 22, wherein the step of comparing the
- 20 current nibble with one or more criteria of the child table is confined to the criteria of one or more criterion groups beginning with the criteria indicated by each of the one or more pointers obtained in a previous identifying step.
24. A multi-stage classification method of selecting a policy from a plurality of policies
- 25 based upon a plurality of protocol data unit (PDU) nibbles, the plurality of policies being embodied in a plurality of hierarchical tables, the plurality of hierarchical tables comprising a first table and one or more child tables, the method comprising the steps of:

- (a) generating, from a first nibble, a first pointer to a first table of a plurality of tables that embody the policies, each table comprising one or more criteria and associated pointer;
 - (b) comparing the first nibble with each of the one or more criteria of the first table beginning with the criterion indicated by the first pointer;
 - (c) identifying one or more matches between the first nibble and one or more associated pointers to a child table;
 - (d) searching a current child tables with a current nibble from the plurality of nibbles, the searching step further comprising:
 - (i) comparing the current nibble with one or more criteria of the child table until a match is detected, beginning with the criteria indicated by each of the one or more pointers obtained in a previous identifying step, and
 - (ii) identifying a pointer to a next child table associated with the match;
 - (e) if no match is detected in the search step:
 - (i) reverting to the previous table, and
 - (ii) resuming the searching step; and
 - (f) repeating the searching step for each child table of the plurality of tables and a policy is matched.
25. The multi-stage classification method of claim 24, wherein the method further includes the steps of determining a score for each match identified in each hierarchical table.
26. The multi-stage classification method of claim 24, wherein the criteria of each hierarchical table are arranged in one or more criterion groups, wherein each criterion of a criterion group in a child hierarchical table shares a common criterion in an associated parent hierarchical table.
27. The multi-stage classification method of claim 27, wherein the step of comparing the current nibble with one or more criteria of the child table is restricted to the criteria of

one or more criterion groups beginning with the criteria indicated by each of the one or more pointers obtained in the previous identifying step.

28. A multi-stage classifier for matching a plurality of strings to one or more policies embodied in a plurality of stages, each stage comprising one or more entries, each entry comprising a criterion and a pointer to a subsequent stage; the classifier being adapted for:

- (a) generating a current pointer to a current stage using a current string;
- (b) searching one or more criteria in the current stage, the searching step further comprising the steps of:
 - (i) comparing the current string to the criteria of one or more entries beginning with entries indicated by the one or more current pointers,
 - (ii) defining the one or more current pointers equal to one or more new pointers, wherein each of the new pointers is associated with a criteria in the current stage that matches the current string, and
 - (iii) incrementing the current stage and the current string;
- (c) repeating the searching step for each of the plurality of stages;
- (d) generating a score for each match identified in each of the plurality of stages; and
- (e) selecting the policy of the one or more policies associated with a highest cumulative score, the cumulative score being the sum of the score for each match at each stage of the plurality of stages.